The role of silent synapses and structural plasticity for memory and amnesia in a model of cortico-hippocampal interplay.

Andreas Knoblauch

2008

Preprint:

This is an accepted article published in Proceedings of the 11th Neural Computation and Psychology Workshop (NCPW). The final authenticated version is available online at: https://doi.org/[DOI not available]

THE ROLE OF STRUCTURAL PLASTICITY AND SYNAPTIC CONSOLIDATION FOR MEMORY AND AMNESIA IN A MODEL OF CORTICO-HIPPOCAMPAL INTERPLAY

A. Knoblauch

Honda Research Institute Europe
Carl-Legien-Strasse 30, D-63073 Offenbach/Main, Germany
E-mail: andreas.knoblauch@honda-ri.de
www.honda-ri.de

This simulation study explores how structural processes and synaptic consolidation during hippocampal memory replay can improve the performance of neocortical neural networks by emulating high effective connectivity in networks that have only low anatomical connectivity. We model ongoing structural plasticity such that, in each time step, a certain fraction of the unconsolidated synapses are eliminated and replaced by new synapses generated at random locations. Simultaneous replay of novel memories consolidates some of the cortical synapses according to Hebbian learning. By this procedure sparsely connected networks can become functionally equivalent to densely connected networks, thereby storing a large amount of information with a tiny number of synapses. In particular, it is possible to store up to $\mathfrak{C}^S \leq \log_2 n$ bits of information per synapse in simple networks of n neurons. This is much more than the well-known bound $\mathfrak{C} \leq 0.72$ bits per synapse for static networks. It turns out that sufficiently fast learning requires a significant number of silent unconsolidated synapses. Thus, with lifetime and stored memories, the number of unconsolidated synapses and thus the ability to learn will decrease gradually. This leads to the discussion of various memory-related effects such as catastrophic forgetting and Ribot gradients in retrograde amnesia.

Keywords: Synaptic plasticity; Associative memory; Willshaw model; Catastrophic forgetting; Ribot gradients.

1. Introduction

Traditionally, learning and memory are attributed to synaptic plasticity, typically by modification of synaptic strength or weight according to variants of the Hebb rule.[1–4] Similarly, artificial neural networks rely almost exclusively on synaptic plasticity in fully connected neural networks.[5] In contrast, connectivity of real neural networks is low even on a local scale. For example, pyramidal cells make synapses to only 10 percent of the neighboring cells within a cubic millimeter of cortical tissue.[6,7] Moreover, plasticity in the brain includes also

structural processes on larger time scales such as elimination and generation of synapses, growth and retraction of spines, and remodeling of dendritic and axonal branches.[8–11] There is increasing evidence that structural plasticity occurs not only during development but is also a regular feature of the adult brain.[12,13]

Here we explore functional implications of structural plasticity[14,15] and synaptic consolidation[16] induced by hippocampal replay[17–19] for storing memories in neural networks of the cerebral cortex. To this end we develop a simple model of structural plasticity and synaptic consolidation and apply it to simple associative networks of the Willshaw-type employing binary synapses.[20–22] It is well known that such network models, in their basic form, can store at most 0.69 bits per synapse, and even more sophisticated models employing real-valued synapses cannot store more than 0.72 bits per synapse.[23–29]

However, our work suggests that, by employing structural plasticity, the storage capacity of these networks could increase to values up to $\log_2 n$ bits per synapse for networks of n neurons.[30–33] This becomes possible because in our model structural plasticity and synaptic consolidation induced by Hebbian learning work together in order to eliminate and replace "useless" synapses by new synapses at possibly more "useful" locations. By this selection procedure a sparsely connected neural network can "place" the rare synapses at the most effective locations and thereby becomes equivalent to a static network with much higher anatomical connectivity. It turns out that sufficiently fast learning consistent with memory transfer from the hippocampus to cortex[18,19,34–37] requires a significant number of silent or unconsolidated synapses which can be re-"placed" by structural plasticity. Thus, with lifetime and stored memories, the number of unconsolidated synapses and, consequently, also the ability to learn will decrease gradually. This leads us finally to the discussion of various memory-related effects such as catastrophic forgetting[38,39] and Ribot gradients in retrograde amnesia.[17–19,34]

2. A simple model of structural plasticity, synaptic consolidation, and cortico-hippocampal interplay

In the following we propose two simple models of structural plasticity that abstract from biological details described in the introduction. For this we apply the concept of a *potential synapse*[15] defined as a cortical location where a presynaptic axon and postsynaptic dendrite are close enough such that a connection could potentially be formed by spine growth and synaptogenesis. We consider the synaptic connections from a neuron population u of size m to another population v of size n that can be described by a synaptic (weight) matrix W of size $m \times n$. We further assume that there are Pmn real synapses and $P_{\rm pot}mn$ potential synapses.

Here P is the *anatomical connectivity* and $P_{\rm pot}$ the *potential connectivity*, where $0 < P < P_{\rm pot}$. Further we assume that each real synapse is either silent (weight 0) or consolidated (weight 1).

For our first model variant (see Fig. 1) we assume that there is at most one potential synapse ij that may connect neuron u_i to neuron v_j. Thus, the network can be described by states $W_{ij} \in \{\nu, \pi, 0, 1\}$. Here state $W_{ij} = \nu$ means that synapse ij does not exist and cannot be realized, $W_{ij} = \pi$ means that ij is a potential synapse not yet realized, $W_{ij} = 0$ means that ij is already realized but still silent, and $W_{ij} = 1$ means that ij is realized and consolidated. States are updated in discrete time steps. Let $p_g := {\rm pr}[W_{ij}(t+1) = 0 | W_{ij}(t) = \nu]$ be the *generation probability* that synapse ij changes from a potential synapse at time t to a real silent synapse at time $t+1$. Similarly, the *elimination probability* is $p_e := {\rm pr}[W_{ij}(t+1) = \nu | W_{ij}(t) = 0]$, the *consolidation probability* is $p_c := {\rm pr}[W_{ij}(t+1) = 1 | W_{ij}(t) = 0]$, and the *deconsolidation probability* is $p_d := {\rm pr}[W_{ij}(t+1) = 0 | W_{ij}(t) = 1]$. In order to keep the total number Pmn of real synapses constant, we have to balance generation and elimination of synapses, for example by choosing $p_g = p_e P_0/(P_{\rm pot} - P)$, where $P_0 mn$ is the number of real silent synapses.

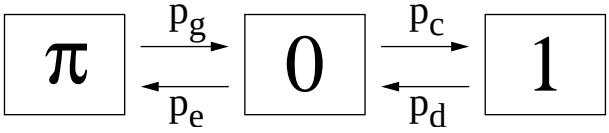

Fig. 1. State diagram illustrating our model of structural plasticity. A synapse can be either potential but not yet realized (state π), realized but still silent (state 0), or realized and consolidated (state 1). Transitions between the states occur with probabilities p_g, p_e, p_c, and p_d as explained in the text.

During the process of *synaptic consolidation* we assume that each synapse receives a binary *consolidation signal* $C_{ij} \in \{0, 1\}$. For a synapse ij we generally assume $p_c = 1$ if $C_{ij} = 1$ and $p_c = 0$ otherwise. Thus, C_{ij} can be interpreted as the "desired" weight matrix for storing a new set of memories in the synaptic connections between neuron populations u and v. Very similar to Hebbian-type learning, the consolidation signal C_{ij} could be provided by coincident presynaptic and postsynaptic neuron activity. The degree to which the actual synaptic connections resemble the "desired" connections can be assessed by the *effective connectivity* $P_{\rm eff} := (\sum_{i,j} C_{ij} W_{ij}) / \sum_{i,j} C_{ij}$ denoting the fraction of "desired" synapses that are actually realized and consolidated. Here the first sum assumes weight zero for states $W_{ij} \in \{\nu, \pi\}$.

To achieve significant *memory consolidation*, C_{ij} must be provided over longer time intervals matching the time scales of structural plasticity. This may be no problem in the case where the new memories correspond to frequently recurring stimuli. However, *episodic memories* typically have only a single training event and thus require a hippocampus-like memory buffer and repetition system that is able to store and replay new memories for some limited time.[34,36] In fact, experiments support the idea that such episodic memories are first buffered by one-shot learning in the hippocampus, and later replayed to become permanently consolidated in the neocortex.[17–19,34–37]

Figure 2 shows a highly simplified model of *cortico-hippocampal interplay* where u and v are interpreted as two cortical neuron populations, and the hippocampus is modeled by an additional neuron population HC. The basic idea is that incoming novel activity patterns are temporarily stored in the connections between neocortex and HC associating each memory u^μ (and v^μ) with an arbitrary index pattern HC$^\mu$. This must happen by one-shot learning such that structural plasticity will be of little use here. However, HC can replay all buffered memories and thereby provide the consolidation signal C_{ij} necessary for final storage of memories in the high capacity cortico-cortical connections from u to v. For this we can assume that the hippocampal indices HC$^\mu$ are organized into a sequence, for example, in the local HC connections such that ordered replay is possible.[40] We can further assume small deconsolidation probability for synaptic connections within neocortex, for example $p_d = 0$, but a much larger deconsolidation probability for connections from, to, and within HC. By this choice, memory lifetime is limited in HC, but virtually unlimited in cortex.

We have also investigated a second model variant of structural plasticity assuming that a newly generated synapse is placed randomly at one of the potential locations. In contrast to the first model, this variant allows multiple realizations of synapses connecting neuron i to neuron j.[41] In fact, if a consolidation signal corresponding to a given set of memories is replayed for too long a time, this will have the effect that all remaining unconsolidated synapses will clutter the few potential locations ij with $C_{ij} = 1$. We have the idea that the first model variant is better suited for spine plasticity where realization of potential synapses is strongly limited by axonal and dendritic geometry (and thus avoids the described cluttering), while the second model variant may be better suited for axonal and dendritic remodeling on a large time scale. In any case, our simulations show that both model variants lead qualitatively to very similar results.[31]

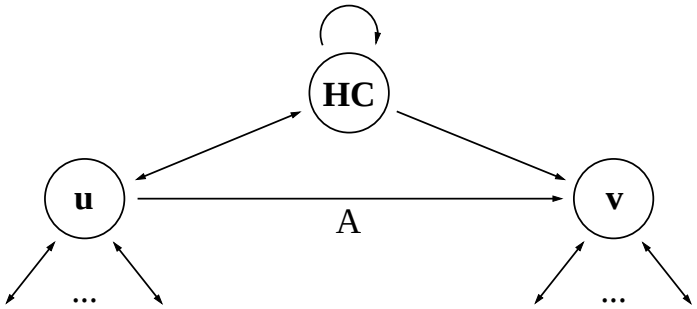

Fig. 2. Model of cortico-hippocampal interplay for memory consolidation. Storing memories means establishing associations between neural activity patterns in cortical areas u and v. These associations are first buffered in the connections from and to the hippocampus (HC) capable of one-shot learning. During the process of consolidation HC can reactivate memories in cortex. By ongoing structural plasticity and synaptic consolidation memories get finally stored in the cortico-cortical connection from u to v. See text for more details.

3. On the function and benefits of structural plasticity for memory storage

We hypothesize that an important function of structural plasticity is to compensate for sparse anatomical connectivity. According to our model, neural networks endowed with structural plasticity are able to increase effective connectivity $P_{\rm eff}$ from the level of anatomical connectivity P towards the level of potential connectivity $P_{\rm pot}$. Thus, such networks may finally become equivalent to static networks with high anatomical connectivity $P \approx P_{\rm pot}$ which would be much more expensive to maintain for the brain in terms of space and energy requirements.[42–44]

Figure 3 illustrates simulations of our first model variant showing that this idea actually works if the number of "required" synapses, $P_{\rm pot} \sum_{i,j} C_{ij}$, does not exceed the number of available synapses, Pmn, or, equivalently, if the *consolidation load* $p_{1C} := (mn)^{-1} \sum_{i,j} C_{ij}$ does not exceed the bound $p_{1C} \leq P/P_{\rm pot}$. The simulations also show that emulating high effective connectivity comes at the price of long replay periods if p_{1C} is close to $P/P_{\rm pot}$. Thus, in order to achieve sufficiently fast consolidation the number of synapses required to be consolidated should be sufficiently smaller (for example factor $1/2$) than the number of available synapses. For more details see[31] where I have given a quantitative analysis of consolidation time required to achieve a desired effective connectivity.

What then are the concrete benefits of structural plasticity and increasing effective connectivity? To answer this question let us consider neural associative networks commonly used as cortical models for storing memories. In our scenario the task is to store a set of M memory associations $u^\mu \rightarrow v^\mu$ ($\mu = 1, \ldots, M$),

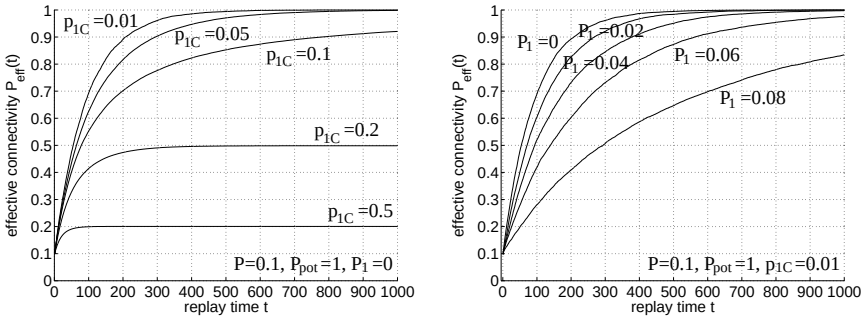

Fig. 3. Results from simulations of the model of structural plasticity for population sizes $m = n = 1000$, anatomical connectivity $P = 0.1$, and full potential connectivity $P_{\rm pot} = 1$. The plots show effective connectivity $P_{\rm eff}(t)$ over replay steps t for different consolidation loads p_{1C} (left panel) and different numbers $P_1 mn$ of initially consolidated synapses (right panel). During each replay time step a fraction $p_e = 0.1$ of the silent synapses was replaced.

where u^μ are address vectors corresponding to activity patterns in the neuron population u having k out of m active units, and, similarly, v^μ are content vectors having l out of n active units. A particularly simple network model is the so-called Willshaw network.[20,21,33,45] In the fully connected static Willshaw network the memories are stored in a binary $m \times n$ weight matrix W by Hebbian learning where $W_{ij} = \min(1, \sum_{\mu=1}^{M} u_i^\mu v_j^\mu)$. This means, a synapse ij has weight 1 iff there is at least one μ with coincident pre- and postsynaptic activity, $u_i^\mu = 1$ and $v_j^\mu = 1$. It turns out that this simple model can store a quite large number of memories, $M \sim mn/\log^2 n$, and that it is possible to store about $\mathfrak{C} \leq 0.7$ bits per binary synapse which is quite close to the information theoretical optimum.[20–22,33] Some theory (e.g., see[22] for a concise description) shows that the storage capacity $\mathfrak{C} \approx {\rm ld} p_{1W} \ln(1 - p_{1W})$ (in bits per synapse) writes as a function of the memory load $p_{1W} := (mn)^{-1} \sum_{ij} W_{ij} = 1 - (1 - kl/(mn))^M$. Here for static networks maximal $\mathfrak{C}$ is achieved for $p_{1W} = 0.5$, while $\mathfrak{C} \to 0$ for low memory load $p_{1W} \to 0$ or large memory load $p_{1W} \to 1$.

However, now consider the Willshaw network endowed with structural plasticity and a consolidation signal identical to the weight matrix of the static network, $C = W$ and $p_{1C} = p_{1W}$. Further we assume, without loss of generality, that the anatomical connectivity is low and equal to the memory load, $P = p_{1W} \ll 1$, and that the potential connectivity is maximal, $P_{\rm eff} = 1$. Because the condition $p_{1C} \leq P/P_{\rm eff}$ is fulfilled the network can emulate full connectivity, $P_{\rm eff} \to 1$, for sufficiently long a consolidation time. This means that the sparsely connected Willshaw network with structural plasticity becomes functionally equivalent to

the fully connected static Willshaw network. Thus, it is possible to store the same number of pattern associations as in the fully connected network, but employing only a small number of synapses. Thus, the storage capacity per synapse is much larger for the network with structural plasticity, $\mathfrak{C}^S = \mathfrak{C}/p_1$. It is easy to see that $\mathfrak{C}^S \to \infty$ for $p_{1W} \to 0$. Thus, in large networks with $n \to \infty$, a single synapse can store an arbitrarily large amount of information. A closer analysis reveals that indeed $\mathfrak{C}^S \sim \log n \to \infty$ (see[30,31,33]). In contrast, it is well known that any static network model of distributed storage cannot exceed the bound $\mathfrak{C} \leq 0.72$ bits per synapse, even if endowed with real-valued synaptic weights.[23–29,46]

Thus, structural plasticity allows us to store large amounts of information with a tiny number of synapses. The intuition is that structural plasticity with hippocampus-like replay and consolidation can "place" the rare synapses at the most useful locations. By this procedure a sparsely connected network can become functionally equivalent to a fully connected network with pruning of irrelevant (i.e., weak or silent) synapses. For that reason, such "zipped" networks can achieve a much higher storage capacity per synapse than static networks.

4. Structural plasticity and catastrophic forgetting

Artificial neural networks are well known for suffering from catastrophic forgetting (CF) also known as the stability-plasticity dilemma.[38] CF means that optimizing synaptic weights for storing a set of new memories will deteriorate or even destroy previous memories. On the other hand, freezing synaptic weights prevents the ability to learn new memories.

In contrast, the learning methods described here for associative networks do not suffer so seriously from CF because the learning contribution for a new memory is independent of other memories. From a functional perspective, associative networks are closely related to look-up tables, where adding a new memory does not affect previous memories.[22] However, associative networks store memories in a distributed way and therefore still may suffer from a weak form of CF, the so-called *Hopfield catastrophe*:[39] This means that a static neural network can store many memories without any problems until the capacity limit is reached. Then storing a single or few further memories can destroy all previously learned memories. This effect is a problem for technical applications but also for modeling memory processes since CF does not normally occur in our brains.

We argue that sparsely connected networks employing structural plasticity do not suffer from CF. Figure 4 (left panel) shows a simulation where we store a larger number of memories exceeding the capacity of the network. Memories are ordered within blocks and each memory block is replayed and consolidated for some time one after the other. The simulations show that approaching the capacity

limit implies that storing new memories becomes more and more difficult, while older memories remain intact even if the capacity limit is exceeded. Thus, in contrast to static networks, there is no CF. The reason is essentially that with storing more and more memories the remaining silent synapses (necessary for storing new memories) become rare and, eventually, are used up right before CF can occur.

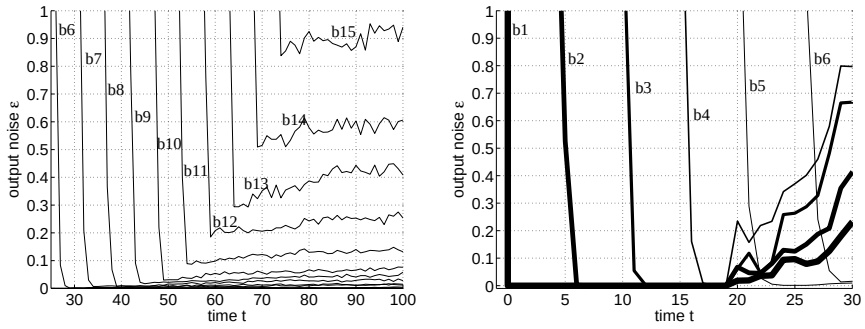

Fig. 4. Results from simulating the model of structural plasticity and cortico-hippocampal interplay for memory consolidation in neocortex. We stored 25 memory blocks (b1,b2,...) each consisting of 4 memories. Each block was replayed by HC for 5 time steps (e.g., block 10 between $t = 45$ and $t = 50$). During each replay step all silent synapses were replaced. The plots show normalized output noise ϵ for each memory block in population v with inactivated HC. Left panel: When storing more and more memories approaching the capacity limit of the network there is a gradual increase of output noise only for new memory blocks while old memories maintain high retrieval quality. Thus, there is no catastrophic forgetting. Right panel: Similar simulation as before, but at time $t = 20$ the cortical network was lesioned by deactivating half of the neurons in population u. This leads to Ribot-like gradients in output noise, i.e., retrieval impairment is more severe for recent memories than remote memories.

5. Retrograde amnesia and Ribot gradients

The same mechanism that prevents CF may be responsible for another salient effect of memory: Patients with lesions of the hippocampus or neighboring neocortex often suffer from graded retrograde amnesia.[34,47–49] This means that lesions impair recent memories more severely than remote memories. These so-called Ribot gradients can also be seen in our model (Fig. 4, right panel). When consolidating more and more memories the number of consolidated synapses (P_1) increases and, correspondingly, the number of unconsolidated silent synapses decreases. Thus, assuming constant replay time per memory block, the effective connectivity that can be achieved for recent memories is smaller than for remote memories (see also Fig. 3, right panel). And this is actually the reason why in our simulations,

after lesions, remote memories are better preserved than recent memories.

In previous theoretical models Ribot gradients have typically been generated by gradients in consolidation time,[17,19] where the Mth memory obtains a $1/M$ share of consolidation time, for example assuming a random walk over the attractor-landscape in Hopfield-type networks with M attractors. Then Ribot gradients occur because early memories can accumulate a much larger total consolidation time (and thus resulting memory trace strength) than recent memories. However, these models implicitly assume that memories are maintained in and consolidated by the hippocampus forever. This contradicts evidence that new memories are buffered by the hippocampus for a limited time only and that replay of novel memories is controlled by the hippocampus.[36,37]

6. Discussion

In this paper we have proposed a simple model of structural plasticity and its relation to synaptic consolidation and cortico-hippocampal interplay. We abstracted from many biological details such as different time scales and geometrical constraints of spine plasticity and remodeling of axons and dendrites. The essence of our model is that structural plasticity can eliminate "useless" synapses (those with low synaptic weights) and regenerate new synapses blindly at potentially more "useful" locations. If a synapse turns out to be actually "useful" it gets consolidated and escapes the process of elimination and regeneration. Since structural plasticity is slow this requires replay of the memories to be consolidated, presumably controlled by the hippocampus.[18,34,35,37]

In contrast to previous approaches[14,15] we apply these ideas to well known associative network models as often used for modeling cortex and memory.[20,21,50,51] By introducing the concept of effective connectivity we have shown that sparsely connected networks with structural plasticity are functionally equivalent to more densely connected static networks. Thus, under some conditions, networks endowed with structural plasticity can store the same large amount of information as fully connected networks, but require only a relatively small number of functional synapses. A closer theoretical analysis reveals that the bits of information stored per synapse can reach the theoretic bound $\log_2 n$ where n is the network size.[30,31,33] Further analyses indicate that these results apply also to biologically more realistic networks based on synapses with gradual weights.[52] In contrast, static neural networks can store at most 0.72 bits per synapse even if endowed with real-valued synapses.[23–29,46] Thus, we propose that the main function of structural plasticity is to emulate higher effective connectivity in networks with sparse anatomical connectivity in order to minimize space and energy requirements.[42–44]

Besides these functional considerations, our model avoids common problems

of static neural networks and can reproduce memory effects found in psychological and neurophysiological experiments. For example, networks endowed with structural plasticity inherently avoid catastrophic forgetting of repeatedly presented memories.[38] Instead, they gradually reduce the capability to acquire new memories, but leave previously stored memories intact. The reason for this behavior is that the number of consolidated synapses will increase with the number of stored memories, and, correspondingly, the number of remaining unconsolidated synapses diminishes. Since silent unconsolidated synapses are necessary for learning new information this process prevents exceeding the storage capacity of the network and thus catastrophic forgetting.

The same mechanism leads to gradients in effective connectivity and thus memory trace strength. Recent memories achieve a lower effective connectivity than remote memories. By this our model can reproduce Ribot gradients as found in patients suffering from retrograde amnesia after cortical lesions.[34,47–49] In previous models[17–19] Ribot gradients have been reproduced by a gradient in total consolidation time. These approaches assume ongoing replay and consolidation of any memory such that the Mth memory gets a time share of only $1/M$. In contrast, our model can reproduce Ribot gradients even for constant replay time per memory. This seems more consistent with common ideas and physiological evidence that new memories get consolidated only for a limited time by hippocampal replay.[18,34–37]

Acknowledgments

The author is grateful to Edgar Körner and Marc-Oliver Gewaltig for providing the opportunity to do this work at the Honda Research Institute. He is also grateful to them and to Ursula Körner, Friedrich Sommer, and Günther Palm for many fruitful discussions and comments.

References

1. D. Hebb, *The organization of behavior. A neuropsychological theory.* (Wiley, New York, 1949).
2. T. Bliss and G. Collingridge, *Nature* **361**, 31 (1993).
3. O. Paulsen and T. Sejnowski, *Current Opinion in Neurobiology* **10**, 172 (2000).
4. S. Song, K. Miller and L. Abbott, *Nature Neuroscience* **3(9)**, 919 (2000).
5. J. Hertz, A. Krogh and R. Palmer, *Introduction to the theory of neural computation.* (Addison-Wesley, Redwood City, 1991).
6. V. Braitenberg and A. Schüz, *Anatomy of the cortex. Statistics and geometry.* (Springer-Verlag, Berlin, 1991).
7. B. Hellwig, *Biological Cybernetics* **82**, 111 (2000).
8. F. Engert and T. Bonhoeffer, *Nature* **399**, 66 (1999).

9. S. Witte, H. Stier and H. Cline, *Journal of Neurobiology* **31**, 219 (1996).
10. C. Woolley, Structural plasticity of dendrites., in *Dendrites.*, eds. G. Stuart, N. Spruston and M. Häusser (Oxford University Press, Oxford, UK, 1999) pp. 339–364.
11. R. Lamprecht and J. LeDoux, *Nature Reviews Neuroscience* **5**, 45 (2004).
12. A. Holtmaat, L. Wilbrecht, G. Knott, E. Welker and K. Svoboda, *Nature* **441**, 979 (2006).
13. V. DePaola, A. Holtmaat, G. Knott, S. Song, L. Wilbrecht, P. Caroni and K. Svoboda, *Neuron* **49**, 861 (2006).
14. P. Poirazi and B. Mel, *Neuron* **29**, 779 (2001).
15. A. Stepanyants, P. Hof and D. Chklovskii, *Neuron* **34**, 275 (2002).
16. S. Fusi, P. Drew and L. Abbott, *Neuron* **45**, 599 (2005).
17. R. Alvarez and L. Squire, *Proceedings of the National Academy of Sciences (USA)* **91**, 7041 (1994).
18. J. McClelland, B. McNaughton and R. O'Reilly, *Psychological Review* **102(3)**, 419 (1995).
19. M. Meeter and J. Murre, *Cognitive Neuropsychology* **22(5)**, 559 (2005).
20. D. Willshaw, O. Buneman and H. Longuet-Higgins, *Nature* **222**, 960 (1969).
21. G. Palm, *Biological Cybernetics* **36**, 19 (1980).
22. A. Knoblauch, *Information Processing Letters* **95**, 537 (2005).
23. M. Tsodyks and M. Feigel'man, *Europhysics Letters* **6**, 101 (1988).
24. G. Palm, On the asymptotic information storage capacity of neural networks., in *Neural Computers*, eds. R. Eckmiller and C. von der Malsburg (Springer Verlag, Berlin, Heidelberg, New York, 1988) pp. 271–280.
25. G. Palm, *Concepts in Neuroscience* **2**, 97 (1991).
26. J.-P. Nadal, *J.Phys. A: Math. Gen.* **24**, 1093 (1991).
27. P. Dayan and D. Willshaw, *Biological Cybernetics* **65**, 253 (1991).
28. G. Palm and F. Sommer, Associative data storage and retrieval in neural nets., in *Models of Neural Networks III*, eds. E. Domany, J. van Hemmen and K. Schulten (Springer-Verlag, New York, 1996) pp. 79–118.
29. A. Knoblauch, *Neural associative networks with optimal Bayesian learning.*, Internal Report HRI-EU 08-xy, Honda Research Institute Europe GmbH (D-63073 Offenbach/Main, Germany, in preparation).
30. A. Knoblauch, *On compressing the memory structures of binary neural associative networks*, Internal Report HRI-EU 06-02, Honda Research Institute Europe GmbH (D-63073 Offenbach/Main, Germany, 2006).
31. A. Knoblauch, *On structural plasticity in neural associative networks.*, Internal Report HRI-EU 08-04, Honda Research Institute Europe GmbH (D-63073 Offenbach/Main, Germany, 2008).
32. A. Knoblauch, *SIAM Journal on Applied Mathematics* **69(1)**, 169 (2008).
33. A. Knoblauch, G. Palm and F. Sommer, *Neural Computation, manuscript under review* (2008).
34. L. Squire and P. Bayley, *Current Opinion in Neurobiology* **17**, 185 (2007).
35. D. Ji and M. Wilson, *Nature Neuroscience* **10(1)**, 100 (2007).
36. R. Ross and H. Eichenbaum, *The Journal of Neuroscience* **26(18)**, 4852 (2006).
37. G. Buzsaki, *Cerebral Cortex* **6**, 81 (1996).
38. R. French, *Trends in Cognitive Sciences* **3(4)**, 128 (1999).

39. A. Robins and S. McCallum, *Connection Science* **7**, 121 (1998).
40. A. Lee and M. Wilson, *Neuron* **36**, 1183 (2002).
41. H. Markram, J. Lübke, M. Frotscher, A. Roth and B. Sakmann, *Journal of Physiology* **500(Pt 2)**, 409 (1997).
42. P. Lennie, *Current Biology* **13**, 493 (2003).
43. S. Laughlin and T. Sejnowski, *Science* **301**, 1870 (2003).
44. D. Attwell and S. Laughlin, *Journal of Cerebral Blood Flow and Metabolism* **21**, 1133 (2001).
45. K. Steinbuch, *Kybernetik* **1**, 36 (1961).
46. E. Gardner and B. Derrida, *J.Phys. A: Math. Gen.* **21**, 271 (1988).
47. S. Zola, Amnesia I: Neuroanatomic and clinical issues., in *Patient-based approaches to cognitive neuroscience.*, eds. M. Farah and T. Feinberg (MIT-Press, Cambridge, MA, 2000) pp. 275–290.
48. A. Baddeley, *Human memory: theory and practice.* (Lawrence Erlbaum, Hillsdale, NJ, 1990).
49. T. Ribot, *Les maladies de la memoire.* (Germer Baillare, Paris, 1881).
50. A. Knoblauch, H. Markert and G. Palm, An associative model of cortical language and action processing., in *Modeling Language, Cognition and Action. Proceedings of the 9th Neural Computation and Psychology, Workshop University of Plymouth, UK, 8-10 September 2004 (NCPW9)*, eds. A. Cangelosi, G. Bugmann and R. Borisyuk (World Scientific Publishing, 2005) pp. 79–86.
51. A. Knoblauch, R. Kupper, M.-O. Gewaltig, U. Körner and E. Körner, *Neurocomputing* **70(10-12)**, 1838 (2007).
52. A. Knoblauch, *Neural associative networks with non-linear local learning.*, Internal Report HRI-EU 09-xy, Honda Research Institute Europe GmbH (D-63073 Offenbach/Main, Germany, in preparation).